arXiv:1606.00996v1 [cs.DS] 3 Jun 2016

A Minimal Variance Estimator for the Cardinality of Big Data Set Intersection

Reuven Cohen Liran Katzir Aviv Yehezkel
Department of Computer Science
Technion
Haifa 32000, Israel

Abstract

In recent years there has been a growing interest in developing "streaming algorithms" for efficient processing and querying of continuous data streams. These algorithms seek to provide accurate results while minimizing the required storage and the processing time, at the price of a small inaccuracy in their output. A fundamental query of interest is the intersection size of two big data streams. This problem arises in many different application areas, such as network monitoring, database systems, data integration and information retrieval. In this paper we develop a new algorithm for this problem, based on the Maximum Likelihood (ML) method. We show that this algorithm outperforms all known schemes and that it asymptotically achieves the optimal variance.

1 Introduction

Classical processing algorithms for database management systems usually require several passes over (static) data sets in order to produce an accurate answer to a user query. However, for a wide range of application domains, the data set is very large and is updated on a continuous basis, making this approach impractical. For this reason, there is a growing interest in developing "streaming algorithms" for efficient processing and querying of continuous data streams in data stream management systems (DSMSs). These algorithms seek to provide accurate results while minimizing both the required storage and the processing time per stream element, at the price of a small inaccuracy in their output [2, 7, 11, 21]. Streaming algorithms for DSMSs typically summarize the data stream using a small sketch, and use probabilistic techniques in order to provide approximate answers to user queries. Such big data streams appear in a wide variety of computer science applications. They are common, for example, in computer networks, where detailed usage statistics (such as the source IP addresses of packets) from different parts of the network need to be continuously collected and analyzed for various security and management tasks.

A fundamental query of interest is the intersection size of two big data streams. Consider two streams of elements, $\overline{A}$ and $\overline{B}$, taken from two sets A and B respectively. Suppose that

each element may appear more than once in each stream. Let $n = \left|\overline{A} \cap \overline{B}\right|$. For example, for $\overline{A} = a, b, c, d, a, b$ and $\overline{B} = a, a, c, c$, we get $n = 2$. Finding n is a problem that arises in many different application areas such as network monitoring, database systems, data integration and information retrieval [4, 5].

As an application example, $a_i \in A$ and $b_j \in B$ could be streams of IP packets passing through two routers, R_1 and R_2. In this case, $|A \cap B|$ represents the number of flows forwarded by both routers. Thus, $\widehat{|A \cap B|}$ can be used for security and traffic monitoring, e.g., on-line detection of denial of service attacks.

As another example, $a_i \in A$ and $b_j \in B$ could be sets of sub-strings found within two text documents. In this case, $|A \cap B|$ represents the number of sub-strings shared by both documents, which can be viewed as the similarity of the documents. Among many other applications, this allows plagiarism detection and pruning of near-duplicate search results in search engines.

One can find the exact value of n by computing the intersection set C in the following way. For every element $b_i \in B$, compare b_i to every $a_j \in A$ and $c_k \in C$. If $b_i \notin C$ and $b_i \in A$, add b_i to C. After all the elements are treated, return the number of elements in C. This naive approach does not scale if storage is limited or if the sets are very large. In these cases, the following estimation problem should be solved. Given two streams of elements (with repetitions) $\overline{A} = a_1, a_2, \ldots, a_p$, and $\overline{B} = b_1, b_2, \ldots, b_q$, such that A and B are the respective sets of the two streams, and $n = |A \cap B|$, find an estimate $\widehat{n}$ of n using only m storage units, where $m \ll n$.

$|A \cap B|$ can be estimated in a straightforward manner using the inclusion-exclusion principle: $\widehat{|A \cap B|} = \widehat{|A|} + \widehat{|B|} - \widehat{|A \cup B|}$, taking advantage of the fact that estimating $|A|$, $|B|$ and $|A \cup B|$ is relatively easy. However, we will later show that this scheme produces inaccurate results. In [2], it is proposed to estimate the Jaccard similarity, defined as $\rho(A, B) = \frac{|A \cap B|}{|A \cup B|}$. The idea is to estimate $|A \cup B|$, and then to extract $|A \cap B|$. A third scheme, proposed in [11], suggests that $\widehat{|A \cap B|} = \frac{\widehat{\rho(A,B)}}{\rho(A,B)+1}(\widehat{|A|} + \widehat{|B|})$. With respect to the above three estimation schemes, the main contributions of this paper are as follows:

1. For the first time, we present a complete analysis of the statistical performance (bias and variance) of the above three schemes.

2. We find the optimal (minimum) variance of any unbiased set intersection estimator.

3. We present and analyze a new unbiased estimator, based on the Maximum Likelihood (ML) method, which outperforms the above three schemes.

The rest of the paper is organized as follows. Section 2 discusses previous work and presents the three previously known schemes. Section 3 presents our new Maximum Likelihood (ML) estimator. It also shows that the new scheme achieves optimal variance and that it outperforms the three known schemes. Section 4 analyzes the statistical performance (bias and variance) of the three known schemes. Section 5 presents simulation results confirming that the new ML estimator outperforms the three known schemes. Finally, Section 6 concludes the paper.

2 Related Work and Previous Schemes

The database research community has extensively explored the problem of data cleaning: detecting and removing errors and inconsistencies from data to improve the quality of databases [23]. Identifying which fields share similar values, identifying join paths, estimating join directions and sizes, and detecting inclusion dependencies are well-studied aspects of this problem [1, 8, 11, 17, 22]. For example, in [11] the authors present several methods for finding related database fields. Their main idea is to hash the values of each field and keep a small sketch that contains the minimal hash values for each. Then, the Jaccard similarity is used to measure similarities between fields. In [1], the authors study the related problem of detecting inclusion dependencies, i.e., pairs of fields A and B such that $A \subseteq B$, and they present an efficient way to test all field pairs in parallel.

All the above problems are closely related to the "cardinality estimation problem" discussed in this paper. This problem has received a great deal of attention in the past decade thanks to the growing number of important real-time "big data" applications, such as estimating the propagation rate of viruses, detecting DDoS attacks [14, 15], and measuring general properties of network traffic [21].

Many works address the cardinality estimation problem [7, 9, 12, 16, 20, 21] and propose statistical algorithms for solving it. These algorithms are usually limited to performing only one pass on the received packets and using a fixed small amount of memory. A common approach is to hash every element into a low-dimensional data sketch, which can be viewed as a uniformly distributed random variable. Then, one of the following schemes is often used to estimate the number of distinct elements in the set:

1. Order-statistics based estimators: In this family of schemes, the identities of the smallest (or largest) k elements are remembered for the considered set. These values are then used for estimating the total number of distinct elements [2, 9, 16, 20]. The family of estimators with $k = 1$ (where the minimal/maximal identity is remembered) is also known as min/max sketches.

2. Bit-pattern based estimators: In this family of schemes, the highest position of the leftmost 1-bit in the binary representation of the identity of each element is remembered and then used for the estimation [7, 12].

If only one hash function is used, the schemes estimate the value of n with an infinite variance. To bound the variance, both schemes repeat the above procedures for m different hash functions and use their combined statistics for the estimation[1].

A comprehensive overview of different cardinality estimation techniques is given in [7, 21]. State-of-the art cardinality estimators have a standard error of about $1/\sqrt{m}$, where m is the number of storage units [6]. The best known cardinality estimator is the HyperLogLog algorithm [12], which belongs to the family of min/max sketches and has a standard error of $1.04/\sqrt{m}$, [12]. For instance, this algorithm estimates the cardinality of a set with 10^9 elements with a standard error of 2% using $m = 2,048$ storage units.

Cardinality estimation algorithms can be used for estimating the cardinality of set intersection. As mentioned in Section 1, a straightforward technique is to estimate the intersection

[1]Stochastic averaging can be used to reduce the number of hash functions from m to only two [13].

using the following inclusion-exclusion principle:

$$\widehat{|A \cap B|} = \widehat{|A|} + \widehat{|B|} - \widehat{|A \cup B|}.$$

This method will be referred to as Scheme-1. Other algorithms first estimate the Jaccard similarity $\rho(A,B) = \frac{|A \cap B|}{|A \cup B|}$, and then use some algebraic manipulation on it [2, 11]. Specifically, the scheme proposed in [2] estimates both the Jaccard similarity and $|A \cup B|$, and then uses

$$\widehat{|A \cap B|} = \widehat{\rho(A,B)} \cdot \widehat{|A \cup B|}.$$

This scheme will be referred to as Scheme-2. The third scheme discussed in this paper, referred to as Scheme-3, is presented in [11]. It estimates the Jaccard similarity, $|A|$, and $|B|$, and then uses

$$\widehat{|A \cap B|} = \frac{\widehat{\rho(A,B)}}{\rho(A,B)+1}(\widehat{|A|} + \widehat{|B|}).$$

The above equation is obtained by substituting the Jaccard similarity definition into $\frac{\rho(A,B)}{\rho(A,B)+1}$, which yields that

$$\frac{\rho(A,B)}{\rho(A,B)+1} = \frac{|A \cap B|}{|A \cap B| + |A \cup B|} = \frac{|A \cap B|}{|A| + |B|}.$$

In [1, 18], the set intersection estimation problem is solved using smaller sample sets. While these techniques are simple and unbiased, they are inaccurate for small sets. In addition, they are sensitive to the arrival order of the data, and to the repetition pattern.

3 A New Maximum Likelihood Scheme with Optimal Variance

In this section we present a new unbiased estimator for the set intersection estimation problem. Because this estimator is based on the Maximum Likelihood (ML) method, it achieves optimal variance and outperforms the three known schemes.

Maximum-Likelihood estimation (ML) is a method for estimating the parameters of a statistical model. For example, suppose we are interested in the height distribution of a given population, but are unable to measure the height of every single person. Assuming that the heights are Gaussian distributed with some unknown mean and variance, the mean and variance can be estimated using ML and only a small sample of the overall population. In general, for a given set of data samples and an underlying statistical model, ML finds the values of the model parameters that maximize the likelihood function, namely, the "agreement" of the selected model with the given sample.

In the new scheme, we first find the (probability density) likelihood function of the set intersection estimation problem, $L(x_A = s, x_B = t; \theta)$; namely, given $\theta = (a, b, n)$ as the

problem parameters[2], this is the probability density of s to be the maximal hash value for A and t to be the maximal hash value for B. Then, we look for values of the problem parameters that maximize the likelihood function.

Table 1 shows some of the notations we use for the rest of the paper.

value	notation
$\|A \cap B\|$	n
$\|A \cup B\|$	u
$\|A\|$	a
$\|B\|$	b
$\|A \setminus B\|$	α
$\|B \setminus A\|$	β

Table 1: Notations

3.1 The Likelihood Function of the Set Intersection Estimation Problem

We first find the likelihood function for one hash function h_k, namely, $L(x_A = s, x_B = t; \theta)_k$, and then generalize it for all hash functions. Recall that for the kth hash function, $x_A^k = \max_{i=1}^{a} \{h_k(a_i)\}$ and $x_B^k = \max_{j=1}^{b} \{h_k(b_j)\}$. To simplify the notation, we shall omit the superscript k for x_A^k and x_B^k, and use x_A and x_B respectively.

We use $\mathrm{PDF}_U(w)$ to denote the probability density function (PDF) of a uniformly distributed random variable $\mathrm{U}(0,1)$ at w. Denote the elements in $A \cap B$ as $\{z_1, z_2, \ldots, z_n\}$. Thus, the elements in A and the elements in B can be written as $\{x_1, x_2, \ldots, x_\alpha, z_1, z_2, \ldots, z_n\}$ and $\{y_1, y_2, \ldots, y_\beta, z_1, z_2, \ldots, z_n\}$ respectively (see Table 1).

We now divide the likelihood function according to the three possible relations between x_A and x_B: $x_A = x_B$, $x_A > x_B$ and $x_A < x_B$.

Case 1: $x_A = x_B$

When $x_A = x_B = s$ holds, the element with the maximal hash value must belong to $A \cap B$. The likelihood function of θ given this outcome is

$$\begin{aligned} L(x_A = x_B = s; \theta) &= \sum_{i=1}^{n} \mathrm{PDF}_U(s) \cdot \Pr\left(x_{(A \cup B) \setminus \{z_i\}} < s\right) \\ &= \sum_{i=1}^{n} s^{u-1} = n \cdot s^{u-1}. \end{aligned} \tag{1}$$

This equality holds because there are n possible elements in $A \cap B$ whose hash value can be the maximum in $A \cup B$, and because $\mathrm{PDF}_U(s) = 1$.

[2]The set intersection estimation problem has six identifying parameters $(n, u, a, b, \alpha, \beta)$, only three of which are needed to derive the others.

Case 2: $x_A < x_B$

In order to have $x_A < x_B$, where $x_A = s$ and $x_B = t$, the maximal hash value in B must also be in $B \setminus A$, and its value must be t. The likelihood function of θ in this case is

$$L(x_{B\setminus A} = t; \theta) = \sum_{j=1}^{\beta} \mathrm{PDF}_U(t) \cdot \Pr\left(x_{(B\setminus A)\setminus\{y_j\}} < t\right) = \beta \cdot t^{\beta-1}.$$

In addition, the maximal hash value in A must be s. The probability density for this is

$$L(x_A = s; \theta) = \sum_{e\in A} \mathrm{PDF}_U(s) \cdot \Pr\left(x_{A\setminus\{e\}} < s\right) = as^{a-1}.$$

Thus,

$$L(x_A < x_B, x_A = s, x_B = t; \theta) = as^{a-1} \cdot \beta t^{\beta-1}. \tag{2}$$

Case 3: $x_A > x_B$

This case is symmetrical to the previous case. Thus, the likelihood function of θ in this case is

$$L(x_A > x_B, x_A = s, x_B = t; \theta) = \alpha s^{\alpha-1} b t^{b-1}.$$

Thus, the likelihood function for set intersection is

$$L(x_A = s, x_B = t; \theta)_k = \begin{cases} ns^{u-1} & x_A = x_B \\ \beta t^{\beta-1} a s^{a-1} & x_A < x_B \\ \alpha s^{\alpha-1} b t^{b-1} & x_A > x_B \end{cases}.$$

We now use the following indicator variables:

1. $I_1 = 1$ if $x_A = x_B$, and $I_1 = 0$ otherwise,
2. $I_2 = 1$ if $x_A < x_B$, and $I_2 = 0$ otherwise,
3. $I_3 = 1$ if $x_A > x_B$, and $I_3 = 0$ otherwise,

to obtain that

$$L(x_A = s, x_B = t; \theta)_k = (ns^{u-1})^{I_1} \cdot (\beta t^{\beta-1} a s^{a-1})^{I_2} \cdot (\alpha s^{\alpha-1} b t^{b-1})^{I_3}. \tag{3}$$

Eq. (3) states the likelihood function for one hash function. To generalize this equation to all m hash functions, denote $S = (s_1, s_2, \ldots, s_m)$ and $T = (t_1, t_2, \ldots, t_m)$ as the m-dimensional vectors of s and t for each hash function.

Corollary 1
The likelihood function for the set intersection estimation problem, for all m hash functions, satisfies:

$$L(x_A = S, x_B = T; \theta) = \prod_{k=1}^{m} (n(s_k)^{u-1})^{I_{1,k}} \cdot (\beta(t_k)^{\beta-1} a(s_k)^{a-1})^{I_{2,k}} \cdot (\alpha(s_k)^{\alpha-1} b(t_k)^{b-1})^{I_{3,k}},$$

where $I_{1,k}$ is the value of I_1 for hash function k, and the same holds for $I_{2,k}$ and $I_{3,k}$. ■

It is usually easier to deal with the log of a likelihood function than with the likelihood function itself. Because the logarithm is a monotonically increasing function, its maximum value is obtained at the same point as the maximum of the function itself. In our case,

$$\begin{aligned}
\log L(x_A = S, x_B = T; \theta) &= \log \prod_{k=1}^{m} L(x_A^k = s_k, x_B^k = t_k; \theta)_k \\
&= \sum_{k=1}^{m} \log L(x_A^k = s_k, x_B^k = t_k; \theta)_k \\
&= \sum_{k=1}^{m} I_{1,k} \cdot \log\left(n \cdot (s_k)^{u-1}\right) + \sum_{k=1}^{m} I_{2,k} \cdot \log\left(\beta \cdot (t_k)^{\beta-1} \cdot a \cdot (s_k)^{a-1}\right) \\
&+ \sum_{k=1}^{m} I_{3,k} \cdot \log\left(\alpha \cdot (s_k)^{\alpha-1} \cdot b \cdot (t_k)^{b-1}\right). \qquad (4)
\end{aligned}$$

3.2 The New Scheme

We use Corollary 1 in order to find $\theta = (a, b, n)$ that maximizes Eq. (4). Let $g(a, b, n)$ and $\mathbb{H}(a, b, n)$ be the gradient and the Hessian matrix of the log-likelihood function. Namely, g is the vector whose components are the partial derivatives of the log-likelihood function for the problem parameters $\theta = (a, b, n)$:

$$g(a, b, n) = \left(\frac{\partial \log L}{\partial a}, \frac{\partial \log L}{\partial b}, \frac{\partial \log L}{\partial n}\right), \qquad (5)$$

and $\mathbb{H}$ is the matrix of the second-order partial derivatives of the log-likelihood function

$$\mathbb{H}_{i,j} = \frac{\partial^2}{\partial \theta_i \partial \theta_j} \log L \quad , 1 \leq i, j \leq 3\,, \qquad (6)$$

where $\theta_1 = a$, $\theta_2 = b$ and $\theta_3 = n$.

The new scheme finds the maximal value of the log-likelihood function, i.e., the root of its gradient $g(a, b, n)$. Practically, this is done using iterations of the Newton-Raphson method on the gradient and Hessian matrix g and H. Starting from an initial estimation $\widehat{\theta_0} = (\widehat{a_0}, \widehat{b_0}, \widehat{n_0})$, the Newton-Raphson method implies that a better estimation is $\widehat{\theta_1} = \widehat{\theta_0} - H^{-1}(\widehat{\theta_0}) \cdot g(\widehat{\theta_0})$. The process is repeated, namely,

$$\widehat{\theta_{l+1}} = \widehat{\theta_l} - H^{-1}(\widehat{\theta_l}) \cdot g(\widehat{\theta_l}), \qquad (7)$$

until a sufficiently accurate estimation is reached. This idea is summarized in the following algorithm.

Algorithm 1
(A Maximum Likelihood scheme for the set intersection estimation problem)
The scheme gets as an input the sketches of the sets $\left\{x_A^k\right\}_{k=1}^m$ *and* $\left\{x_B^k\right\}_{k=1}^m$*, where* x_A^k *and* x_B^k *are the maximal hash values of* A *and* B *respectively for the* k*th hash function, and returns an estimate of their set intersection cardinality.*

1) *Estimate* $a_0 = \widehat{a}$, $b_0 = \widehat{b}$ *and* $\widehat{u}$ *using any cardinality estimation algorithm, such as [12].*

2) *Estimate the Jaccard similarity* $\widehat{\rho}$ *from the given sketches of* A *and* B*.*

3) *Find the maximum of the likelihood function* L *(Eq. (4)) as explained above; use* $n_0 = \widehat{\rho} \cdot \widehat{u}$ *as an initial value of* n *(see Scheme-2 in Section 1), and* a_0, b_0 *as an initial values of* a *and* b *respectively.*

4) *Return* $\widehat{n}$*.*

When we implemented Algorithm 1, we discovered that 3 Newton-Raphson iterations are enough for the algorithm to converge.

3.3 The Optimal Variance of the New Estimator

The new estimator proposed in this section is based on Maximum Likelihood and thus it asymptotically achieves optimal variance [24]. We use the Cramer-Rao bound to compute this optimal variance.

The Cramer-Rao bound states that the inverse of the Fisher information matrix is a lower bound on the variance of any unbiased estimator of θ [24]. The Fisher information matrix $\mathbb{F}_{i,j}$ is a way of measuring the amount of information that a random variable X carries about an unknown parameter θ upon which the probability of X depends. It is defined as:

$$\mathbb{F}_{i,j} = -\mathbb{E}\left[\frac{\partial^2}{\partial\theta_i \partial\theta_j} \log L\right]. \tag{8}$$

We now use the log-likelihood function (Eq. (4)) to derive this matrix for the set intersection estimation problem:

$$\mathbb{F}(a,b,n) = m \cdot \begin{pmatrix} \frac{\beta}{u}\cdot\frac{1}{a^2} + \frac{1}{u\cdot\alpha} & 0 & \frac{-1}{u\cdot\alpha} \\ 0 & \frac{\alpha}{u}\cdot\frac{1}{b^2} + \frac{1}{u\cdot\beta} & \frac{-1}{u\cdot\beta} \\ \frac{-1}{u\cdot\alpha} & \frac{-1}{u\cdot\beta} & \frac{1}{u\cdot n} + \frac{1}{u\cdot\beta} + \frac{1}{u\cdot\alpha} \end{pmatrix}, \tag{9}$$

where each term is derived due to algebraic manipulations and derivatives of the log-likelihood function. Note that the expected values of the indicator variables $I_{1,k}, I_{2,k}$ and $I_{3,k}$ are required to derive the matrix. For $I_{1,k}$ we get:

$$\mathbb{E}[I_{1,k}] = \Pr\left(x_A^k = x_B^k\right) = \frac{n}{u} \quad , 1 \le k \le m\,.$$

The first equality is due to the definition of $I_{1,k}$, and the second is due to Eq. (11). The following are obtained in the same way for every $1 \le k \le m$:

1. $\mathbb{E}[I_{2,k}] = \Pr\left(x_A^k < x_B^k\right) = \frac{\beta}{u}$.
2. $\mathbb{E}[I_{3,k}] = \Pr\left(x_A^k > x_B^k\right) = \frac{\alpha}{u}$.

Let $\widehat{|A \cap B|}$ be an unbiased estimator for the set intersection estimation problem. Then, according to the Cramer-Rao bound, $\text{Var}\left[\widehat{|A \cap B|}\right] \geq (\mathbb{F}^{-1})_{3,3}$, where $(\mathbb{F}^{-1})_{3,3}$ is the term in place $[3,3]$ in the inverse Fisher information matrix. Finally, from the computation of the term $(\mathbb{F}^{-1})_{3,3}$, we can obtain the following corollary:

Corollary 2
$\text{Var}\left[\widehat{|A \cap B|}\right] \geq (\mathbb{F}^{-1})_{3,3}$, *where,*

$(\mathbb{F}^{-1})_{3,3} = \frac{n \cdot u}{m} \cdot \frac{(b^2+\alpha\beta)(a^2+\alpha\beta)}{\alpha \cdot n(a^2+\alpha\beta)+\beta \cdot n(b^2+\alpha\beta)+(a^2+\alpha\beta)(b^2+\alpha\beta)}$.

4 An Analysis of the Three Schemes From Section 1

In this section we will analyze the statistical performance (bias and variance) of the three schemes discussed in Section 1 for set intersection estimation.

4.1 Preliminaries

4.1.1 Jaccard Similarity

Recall that the Jaccard similarity is defined as: $\rho(A,B) = \frac{|A \cap B|}{|A \cup B|}$, where A and B are two finite sets. Its value ranges between 0, when the two sets are completely different, and 1, when the sets are identical. An efficient and accurate estimate of ρ can be computed as follows [3]. Each item in A and B is hashed into $(0,1)$, and the maximal value of each set is considered as a sketch that represents the whole set. To improve accuracy, m hash functions are used[3], and the sketch of each set is a vector of m maximal values. Given a set $A = \{a_1, a_2, \ldots, a_p\}$ and m different hash functions $h_1, h_2, \ldots, h_m$, the maximal hash value for the jth hash function can be formally expressed as:

$$x_A^j = \max_{i=1}^{p} \{h_j(a_i)\} \quad , \quad 1 \leq j \leq m,$$

and the sketch of A is:

$$X_{(A)} = \left\{x_A^1, x_A^2, \ldots, x_A^m\right\}.$$

$X_{(B)}$ is computed in the same way. Then, the two sketches are used to estimate the Jaccard similarity of A and B:

$$\widehat{\rho(A,B)} = \frac{\sum_{j=1}^{m} I_{x_A^j == x_B^j}}{m}, \tag{10}$$

where the indicator function $I_{x_A^j == x_B^j}$ is 1 if $x_A^j = x_B^j$, and 0 otherwise. To shorten our notation, for the rest of the paper we use ρ to indicate $\rho(A,B)$.

[3]Stochastic averaging can be used to reduce the number of hash functions from m to only two [13].

Lemma 1
In Eq. (10), $\widehat{\rho}$ is a normally distributed random variable with mean ρ and variance $\frac{1}{m}\rho(1-\rho)$; i.e., $\widehat{\rho} \to \mathcal{N}\left(\rho, \frac{1}{m}\rho(1-\rho)\right)$.

Proof:
Consider the j-th hash function. According to [3],

$$\Pr\left(x_A^j = x_B^j\right) = \frac{|A \cap B|}{|A \cup B|}. \tag{11}$$

The intuition is to consider the hash function h_j and define $m(S)$, for every set S, to be the element in S with the maximum hash value of h_j, i.e., $h_j(m(S)) = x_S^j$. Then, we get $m(A) = m(B)$ only when $m(A \cup B)$ lies also in their intersection $A \cap B$. The probability of this is the Jaccard ratio ρ, and therefore $\Pr\left(x_A^j = x_B^j\right) = \rho$.

From Eqs. (10) and (11) follows that $\widehat{\rho}$ is a sum of m Bernoulli variables. Therefore, it is binomially distributed, and can be asymptotically approximated to normal distribution as $m \to \infty$; namely, $\widehat{\rho} = \frac{\sum_{l=1}^{m} I_{x_A^j = x_B^j}}{m} \to \mathcal{N}\left(\rho, \frac{1}{m}\rho(1-\rho)\right)$.

4.1.2 The Cardinality Estimation Problem

Algorithms for estimating the cardinality of set intersection use estimations of $|A|$, $|B|$, and $|A \cup B|$. These estimations can be found using well-known algorithms for the following cardinality estimation problem:

Instance: A stream of elements $x_1, x_2, \ldots, x_s$ with repetitions. Let c be the number of different elements, namely $c = |\{x_1, x_2, \ldots, x_s\}|$.

Objective: Find an estimate $\widehat{c}$ of c using only m storage units, where $m \ll c$.

For the rest of the paper we consider the HyperLogLog algorithm [12] for solving the above problem. As indicated in Section 2, this algorithm has a very small standard error, of about $1.04/\sqrt{m}$ where m is the number of storage units. The pseudo-code of this algorithm is as follows:

Algorithm 2 *The HyperLogLog algorithm for the cardinality estimation problem*

1. *Initialize m registers: $C_1, C_2, \ldots, C_m$ to 0.*
2. *For each input element x_i do:*
   (a) *Let $\rho = \lfloor -\log_2(h_1(x_i)) \rfloor$ be the leftmost 1-bit position of the hashed value.*
   (b) *Let $j = h_2(x_i)$ be the bucket for this element.*
   (c) *$C_j \leftarrow \max\{C_j, \rho\}$.*
3. *To estimate the value of n do:*
   (a) *$Z \leftarrow \left(\sum_{j=1}^{m} 2^{-C_j}\right)^{-1}$ is the harmonic mean of 2^{C_j}.*

(b) return $\alpha_m m^2 Z$, *where*
$\alpha_m = \left(m \int_0^\infty \left(\log_2 \left(\frac{2+u}{1+u}\right)\right)^m \, du\right)^{-1}$.

The following lemma summarizes the statistical performance of Algorithm 2:

Lemma 2
For Algorithm 2, $\widehat{c} \to \mathcal{N}\left(c, \frac{c^2}{m}\right)$, *where* c *is the actual cardinality of the considered set,* $\widehat{c}$ *is the estimate computed by the algorithm, and* m *is the number of storage units used by the algorithm. When Algorithm 2 is used with two sets* A *and* B, *the following holds:*

$$\widehat{|A|} \to \mathcal{N}\left(|A|, \frac{|A|^2}{m}\right),$$

$$\widehat{|B|} \to \mathcal{N}\left(|B|, \frac{|B|^2}{m}\right),$$

$$\textit{and} \quad \widehat{|A \cup B|} \to \mathcal{N}\left(|A \cup B|, \frac{|A \cup B|^2}{m}\right). \tag{12}$$

The proof is given in [12].

Let us also recall three general lemmas, not related to set intersection cardinality estimation. The first lemma, known as the Delta Method, allows us to compute the probability distribution for a function of an asymptotically normal estimator using the estimator's variance:

Lemma 3 (Delta Method)
Let $\theta_1, \theta_2, \ldots, \theta_m$ *be a sequence of* m *random variables such that for every integer* i $\sqrt{i}(\theta_i - \theta) \to \mathcal{N}(0, \sigma^2)$, *where* θ *and* σ^2 *are finite valued constants. Then, for every integer* i *and for every function* g *for which* $g'(\theta)$ *exists and* $g'(\theta) \neq 0$, *the following holds:*

$$\sqrt{i}(g(\theta_i) - g(\theta)) \to \mathcal{N}\left(0, \sigma^2 g'(\theta)^2\right).$$

A proof is given in [24].

The next lemma shows how to compute the probability distribution of a random variable that is a product of two normally distributed random variables whose covariance is 0:

Lemma 4 (Product distribution)
Let X *and* Y *be two random variables satisfying* $X \to \mathcal{N}(\mu_x, \sigma_x^2)$ *and* $Y \to \mathcal{N}\left(\mu_y, \sigma_y^2\right)$, *such that* $\mathrm{Cov}[X, Y] = 0$. *Then, the product* $X \cdot Y$ *asymptotically satisfies the following:*

$$X \cdot Y \to \mathcal{N}\left(\mu_x \mu_y, \mu_y^2 \sigma_x^2 + \mu_x^2 \sigma_y^2\right).$$

A proof is given in [24].

The final lemma states the distribution of the maximal hash value. Let us first recall the beta distribution. $\mathrm{Beta}(\alpha, \beta)$ is defined over the interval $(0, 1)$ and has the following probability and cumulative density functions (PDF and CDF respectively):

$$f(x) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\beta)\Gamma(\alpha)} x^{\alpha-1}(1-x)^{\beta-1}$$

$$F(x) = \int_0^x \frac{\Gamma(\alpha+\beta)}{\Gamma(\beta)\Gamma(\alpha)} x^{\alpha-1}(1-x)^{\beta-1} dx,$$

where $\Gamma(z)$ is the gamma function, defined as $\int_0^\infty e^{-t}t^{z-1}dt$. Using integration by parts, the gamma function can be shown to satisfy $\Gamma(z+1) = z \cdot \Gamma(z)$. Combining this with $\Gamma(1) = 1$ yields that $\Gamma(n) = (n-1)!$ holds for every integer n. Two other known beta identities are [19]:

$$\mathrm{Beta}(1, 1) \sim \mathrm{U}(0, 1)$$

and

$$\mathrm{Beta}(\alpha, \beta) \sim 1 - \mathrm{Beta}(\beta, \alpha).$$

The following lemma presents some key properties of the beta distribution, which we will use in the analysis.

Lemma 5
Let $x_1, x_2, \ldots, x_n$ be independent RVs, where $x_i \sim \mathrm{U}(0, 1)$. Then,

(a) $X = \max_{i=1}^n x_i \sim \mathrm{Beta}(n, 1)$.

(b) X *satisfies the following*

(1) $\mathbb{E}[X] = \frac{n}{n+1}$; *and*
(2) $\mathrm{Var}[X] = \frac{n}{(n+1)^2(n+2)}$.

A proof for (a) is given in [10]; the other equalities follow the beta distribution of X.

4.2 Analysis of Scheme-1

Scheme-1 estimates the cardinality of $A \cap B$ using the inclusion-exclusion principle:

$$\widehat{|A \cap B|} = \widehat{|A|} + \widehat{|B|} - \widehat{|A \cup B|}.$$

Let $\widehat{n_1}$ be the estimator found by Scheme-1. The following theorem summarizes its statistical performance.

Theorem 3
$\frac{\widehat{n_1}}{n} \to \mathcal{N}\left(1, \frac{1}{mn^2}(u^2 - a^2 - b^2) - \frac{2a \cdot b}{m \cdot u \cdot n} + \frac{2u \cdot (a^2(b^2+\alpha\beta)+b^2(a^2+\alpha\beta))}{m \cdot Z \cdot n}\right)$, *where m is the number of storage units, and Z satisfies:*

$$Z = \alpha \cdot n(a^2 + \alpha\beta) + \beta \cdot n(b^2 + \alpha\beta) + (a^2 + \alpha\beta)(b^2 + \alpha\beta).$$

Proof:
For the expectation, from Lemma 2 follows that:

$$\mathbb{E}[\widehat{n_1}] = \mathbb{E}[\widehat{a}] + \mathbb{E}\left[\widehat{b}\right] - \mathbb{E}[\widehat{u}] = \\ = a + b - u = n.$$

The first equality is due to the definition of Scheme-1 and the expectation properties, and the second equality is due to Lemma 2. Thus, the estimator is unbiased. For the variance, Lemma 7 in the Appendix proves that $\mathrm{Cov}[a,b] = \frac{n \cdot a \cdot b}{m \cdot u}$, and Lemma 8 in the Appendix proves that $\mathrm{Cov}[a,u] = \frac{a^2}{m} + \frac{n \cdot a \cdot b}{m \cdot u} - \frac{u \cdot n \cdot a^2(b^2+\alpha\beta)}{m \cdot Z}$ and $\mathrm{Cov}[b,u] = \frac{b^2}{m} + \frac{n \cdot a \cdot b}{m \cdot u} - \frac{u \cdot n \cdot b^2(a^2+\alpha\beta)}{m \cdot Z}$, where $Z = \alpha \cdot n(a^2+\alpha\beta) + \beta \cdot n(b^2+\alpha\beta) + (a^2+\alpha\beta)(b^2+\alpha\beta)$. We get that:

$$\begin{aligned}
\mathrm{Var}[\widehat{n_1}] &= \mathrm{Var}[\widehat{a}] + \mathrm{Var}\left[\widehat{b}\right] + \mathrm{Var}[\widehat{u}] + 2\,\mathrm{Cov}\left[\widehat{a},\widehat{b}\right] - 2\,\mathrm{Cov}[\widehat{a},\widehat{u}] - 2\,\mathrm{Cov}\left[\widehat{b},\widehat{u}\right] = \\
&= \frac{a^2}{m} + \frac{b^2}{m} + \frac{u^2}{m} + 2 \cdot \frac{n \cdot a \cdot b}{m \cdot u} \\
&- 2 \cdot \frac{a^2}{m} - 2 \cdot \frac{n \cdot a \cdot b}{m \cdot u} + 2 \cdot \frac{u \cdot n \cdot a^2(b^2+\alpha\beta)}{m \cdot Z} \\
&- 2 \cdot \frac{b^2}{m} - 2 \cdot \frac{n \cdot a \cdot b}{m \cdot u} + 2 \cdot \frac{u \cdot n \cdot b^2(a^2+\alpha\beta)}{m \cdot Z} = \\
&= \frac{1}{m}(u^2 - a^2 - b^2) - \frac{2n \cdot a \cdot b}{m \cdot u} + \frac{2u \cdot n \cdot (a^2(b^2+\alpha\beta) + b^2(a^2+\alpha\beta))}{m \cdot Z}. \qquad (13)
\end{aligned}$$

The first equality is due to variance properties and because a, b and u are dependent ($\mathrm{Cov}[a,b]$, $\mathrm{Cov}[a,u]$ and $\mathrm{Cov}[b,u]$ are all $\neq 0$). The second equality is due to Lemma 2, Lemma 7 and Lemma 8 (both are in the Appendix). The third equality is due to algebraic manipulations. Finally, after dividing by n^2 we get the result.

4.3 Analysis of Scheme-2

Scheme-2 estimates the cardinality of $A \cap B$ by estimating the Jaccard similarity ρ and $|A \cup B|$ [2]:

$$\widehat{|A \cap B|} = \widehat{\rho} \cdot \widehat{|A \cup B|}.$$

Let $\widehat{n_2}$ be the estimator found by Scheme-2. The following theorem summarizes its statistical performance.

Theorem 4
$\frac{\widehat{n_2}}{n} \to \mathcal{N}\left(1, \frac{1}{m\rho}\right)$, *where m is the number of storage units.*

Proof:
From the definition of Scheme-2, $\widehat{n_2} = \widehat{\rho} \cdot \widehat{u}$. From Lemma 1 and Lemma 2 follows that:

1. $\widehat{\rho} \to \mathcal{N}\left(\rho, \frac{1}{m}\rho(1-\rho)\right)$.

2. $\widehat{u} \to \mathcal{N}\left(u, \frac{u^2}{m}\right)$.

Applying Lemma 4 for the expectation yields:

$$\mathbb{E}[\widehat{n_2}] = \mathbb{E}[\widehat{\rho} \cdot \widehat{u}] = \rho \cdot u = n.$$

Therefore, the estimator is unbiased. For the variance, applying again Lemma 4 yields:

$$\begin{aligned}\mathrm{Var}[\widehat{n_2}] &= u^2 \cdot \frac{1}{m}\rho(1-\rho) + \rho^2 \cdot \frac{u^2}{m} \\ &= \frac{1}{m}(u \cdot n - n^2 + n^2) \\ &= \frac{n^2}{m\rho},\end{aligned}$$

where all the equalities are due to Lemma 4 and algebraic manipulations. Finally, after dividing by n^2, we get $\frac{\widehat{n_2}}{n} \to \mathcal{N}\left(1, \frac{1}{m\rho}\right)$.

4.4 Analysis of Scheme-3

Scheme-3 estimates the cardinality of $A \cap B$ by estimating the Jaccard similarity ρ, $|A|$, and $|B|$ [11]:

$$\widehat{|A \cap B|} = \frac{\widehat{\rho}}{\rho + 1}(\widehat{|A|} + \widehat{|B|}). \tag{14}$$

Let $\widehat{n_3}$ be the estimator found by Scheme-3. We will use the following lemma in the analysis.

Lemma 6
$\widehat{a} + \widehat{b} \to \mathcal{N}\left(a + b, \frac{1}{m}(a^2 + b^2 + 2ab\rho)\right)$, *where m is the number of storage units.*

Proof:
For the expectation,

$$\mathbb{E}\left[\widehat{a} + \widehat{b}\right] = \mathbb{E}[\widehat{a}] + \mathbb{E}\left[\widehat{b}\right] = a + b. \tag{15}$$

The first equality is due to expectation properties, and the second is due to Lemma 2. For the variance, Lemma 7 in the Appendix proves that $\mathrm{Cov}[a, b] = \frac{n \cdot a \cdot b}{m \cdot u}$. Thus,

$$\begin{aligned}\mathrm{Var}\left[\widehat{a} + \widehat{b}\right] &= \\ &= \mathrm{Var}[\widehat{a}] + \mathrm{Var}\left[\widehat{b}\right] + 2\,\mathrm{Cov}\left[\widehat{a}, \widehat{b}\right] = \\ &= \frac{a^2}{m} + \frac{b^2}{m} + 2 \cdot \frac{n \cdot a \cdot b}{m \cdot u} = \\ &= \frac{1}{m}(a^2 + b^2 + 2ab\rho).\end{aligned} \tag{16}$$

The first equality is due to variance properties and because a and b are dependent ($\mathrm{Cov}[a,b] \neq 0$). The second equality is due to Lemma 2 and Lemma 7, and the third equality is due to algebraic manipulations and the Jaccard similarity definition ($\rho = \frac{n}{u}$). Combining Eqs. 15 and 16 yields that

$$\widehat{a} + \widehat{b} \to \mathcal{N}\left(a+b, \frac{1}{m}(a^2+b^2+2ab\rho)\right). \tag{17}$$

The following theorem summarizes the statistical performance of $\widehat{n_3}$.

Theorem 5
$\frac{\widehat{n_3}}{n} \to \mathcal{N}\left(1, \frac{1}{m}(1+\frac{2ab}{u(a+b)}+(\alpha+\beta)\frac{u^2}{n(a+b)^2})\right)$, *where m is the number of storage units.*

Proof:
From the definition of Scheme-3, $\widehat{n_3} = \widehat{\frac{\rho}{\rho+1}}(\widehat{a}+\widehat{b})$. From Lemma 1 follows that:

$$\widehat{\rho} \to \mathcal{N}\left(\rho, \frac{1}{m}\rho(1-\rho)\right).$$

Applying Lemma 3 on $\frac{\rho}{\rho+1}$ yields:

$$\widehat{\frac{\rho}{\rho+1}} \to \mathcal{N}\left(\frac{\rho}{\rho+1}, \frac{1}{m}\rho(1-\rho)\frac{1}{(1+\rho)^4}\right). \tag{18}$$

Applying Lemma 4 for Eqs. (18) and (17) yields:

$$\mathbb{E}[\widehat{n_3}] = \mathbb{E}\left[\widehat{\frac{\rho}{\rho+1}} \cdot \widehat{a+b}\right] = \frac{\rho}{\rho+1}\cdot(a+b) = n.$$

Therefore, the estimator is unbiased. For the variance, applying again Lemma 4 yields:

$$\mathrm{Var}[\widehat{n_3}] = \frac{1}{m}(n^2 + \frac{2abn^2}{u(a+b)} + (\alpha+\beta)\frac{u^2 n}{(a+b)^2}).$$

Finally, after dividing by n^2, we get: $\frac{\widehat{n_3}}{n} \to \mathcal{N}\left(1, \frac{1}{m}(1+\frac{2ab}{u(a+b)}+(\alpha+\beta)\frac{u^2}{n(a+b)^2})\right)$.

A simple comparison yields that $\mathrm{Var}[\widehat{n_2}] > \mathrm{Var}[\widehat{n_3}]$, i.e., Scheme-3 outperforms Scheme-2. However, a similar comparison between Scheme-1 and Scheme-3 cannot be done, because neither one is always better than the other.

5 Simulation Results

In this section we examine the performance of our new ML estimator and show that it indeed outperforms the three known schemes. We implemented all four schemes, and simulated two sets, A and B, whose cardinalities are as follows:

1. $|A| = a = 10^6$;

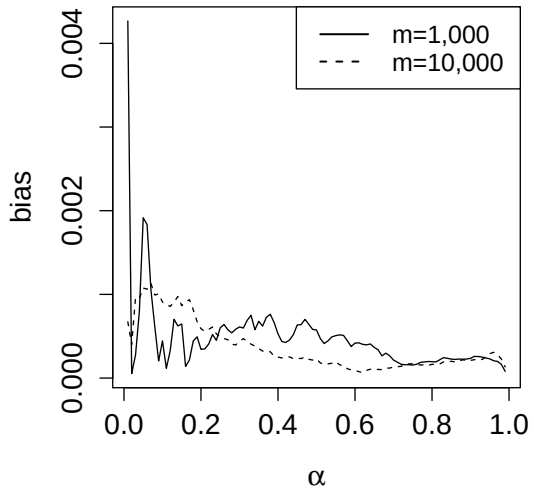

Figure 1: The bias of the new ML scheme for $f = 1$ and for different α values, for $m = 1,000$ and $m = 10,000$

2. $|B| = a \cdot f$, where $f > 0$;
3. $|A \cap B| = a \cdot \alpha$, where $0 \leq \alpha \leq 1$.

We estimate $\widehat{|A \cap B|}$ for each of the four schemes, for $f \in \{1, 5, 10\}$, and for $\alpha \in \{0, 0.01, 0.02, \ldots, 0.98, 0.99, 1\}$. We repeat the test for $10,000$ different sets. Thus, for each of the four schemes, and for each f and α values, we get a vector of $10,000$ different estimations. Then, for each f and α values, we compute the variance and bias of this vector, and view the result as the variance and bias of the estimator (for the specific f and α values). Each such computation is represented by one point in the graph. Let $v_{f,\alpha} = (\widehat{n}_1, \ldots, \widehat{n}_{10^4})$ be the vector of estimations for a specific scheme and for specific f and α values. Let $\mu = \frac{1}{10^4} \sum_{i=1}^{10^4} \widehat{n}_i$, be the mean of $v_{f,\alpha}$. The bias and variance of $v_{f,\alpha}$ are computed as follows:

$$\text{Bias}(v_{f,\alpha}) = \left| \frac{1}{n}(\mu - n) \right|$$

and

$$\text{Var}\left[v_{f,\alpha}\right] = \frac{1}{10^4} \sum_{i=1}^{10^4} (\widehat{n}_i - \mu)^2.$$

Figure 1 presents the bias of the ML estimator for $f = 1$, different α values, and two values of m: $m = 1,000$ and $m = 10,000$ (recall that m is the number of hash values used for the estimations of $\widehat{|A|}$, $\widehat{|B|}$ and the Jaccard similarity $\widehat{\rho}$). We can see that the bias is very small for all α and m values. We got very similar results for bigger f values as well.

Figure 2 presents the normalized variance ($\text{Var}\left[\frac{\widehat{n}}{n}\right]$) of our ML estimator for different f and α values, and for $m \in \{100, 500, 1000, 10000\}$. As expected, the normalized variance

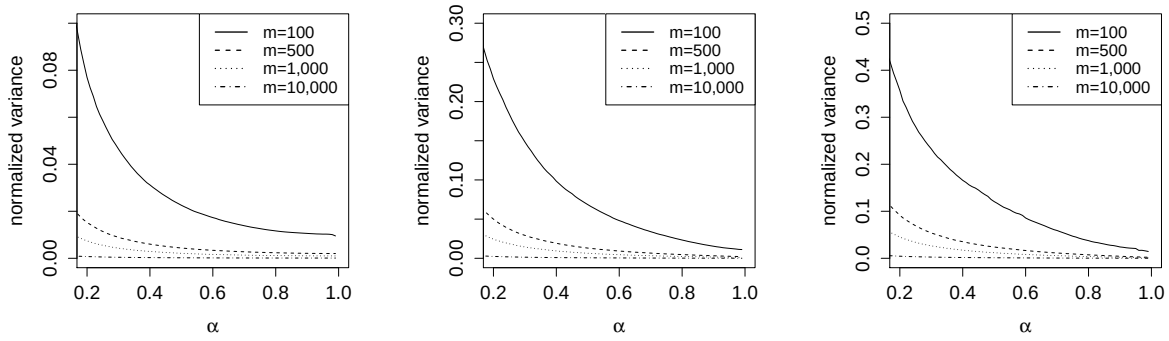

(a) $f = 1$ (b) $f = 5$ (c) $f = 10$

Figure 2: The normalized variance of the new ML scheme for different values of f, α and m

decreases as the number of hash values (m) increases, or as α increases. Overall, the normalized variance is very small for all values of α, f and m, indicating that the new scheme is very precise.

After showing that the new ML scheme indeed yields good results, we now compare its performance to that of Schemes 1-3. When comparing the statistical performance of two algorithms, it is common to look at their MSE (mean squared error) or RMSE, where $\text{MSE} = (\text{Bias}(\widehat{\theta}))^2 + \text{Var}\left[\widehat{\theta}\right]$, and $\text{RMSE} = \sqrt{(\text{Bias}(\widehat{\theta}))^2 + \text{Var}\left[\widehat{\theta}\right]}$. In our case, because all the estimators are unbiased, we compare only their variance. We define the "relative variance improvement" of the new ML scheme over each scheme as

$$\frac{\text{Var}\left[\widehat{\theta_i}\right] - \text{Var}\left[\widehat{\theta_{ML}}\right]}{\text{Var}\left[\widehat{\theta_i}\right]},$$

where θ_i is the estimator of the ith scheme, and θ_{ML} is the new ML estimator.

Figure 3 presents the simulation results for two values of m: $m = 10,000$ (upper graphs) and $m = 1,000$ (lower graphs). **We can see that the new ML estimator outperforms the three schemes for all values of α and f, and for both values of m.** This improvement varies between 100% to a few percent.

6 Conclusion

In this paper we studied the problem of estimating the number of distinct elements in the set intersection of two streams. We presented a complete analysis of the statistical performance (bias and variance) of three previously known schemes. We then computed the likelihood function of the problem and used it to present a new estimator, based on the ML method. We also found the optimal variance of any unbiased set intersection estimator, which is asymptotically achieved by our new ML scheme. We can conclude that our new scheme outperforms the three known schemes, significantly improves the variance (precision) of the estimator, and yields better results than the three previously known schemes.

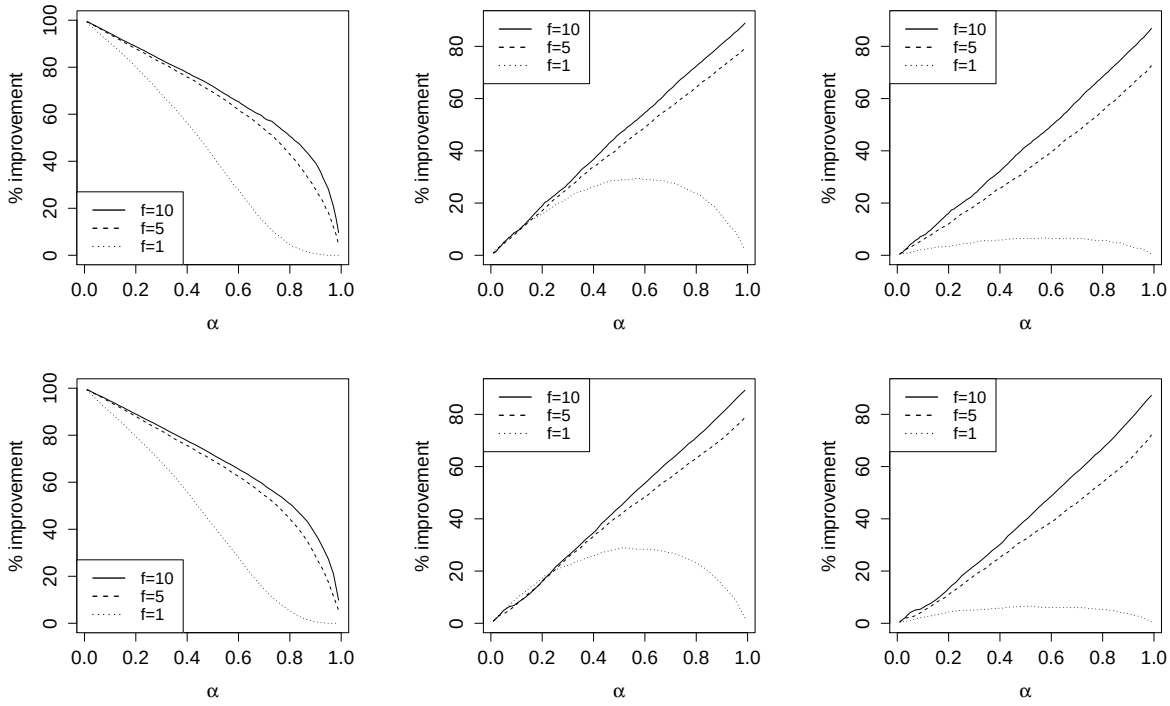

(a) Scheme-1 (b) Scheme-2 (c) Scheme-3

Figure 3: The percentage of variance improvement of the new ML scheme over each of the other schemes for different values of f and α; the upper graphs are for $m = 10,000$ and the lower graphs are for $m = 1,000$

References

[1] J. Bauckmann, U. Leser, F. Naumann, and V. Tietz. Efficiently detecting inclusion dependencies. In *ICDE*, pages 1448–1450, 2007.

[2] K. S. Beyer, P. J. Haas, B. Reinwald, Y. Sismanis, and R. Gemulla. On synopses for distinct-value estimation under multiset operations. In *SIGMOD Conference*, pages 199–210, 2007.

[3] A. Z. Broder. On the resemblance and containment of documents. In *IEEE Compression and Complexity of Sequences 1997*, pages 21–29.

[4] A. Z. Broder. Identifying and filtering near-duplicate documents. In *CPM*, pages 1–10, 2000.

[5] A. Z. Broder, S. C. Glassman, M. S. Manasse, and G. Zweig. Syntactic clustering of the web. *Computer Networks*, 29(8-13):1157–1166, 1997.

[6] P. Chassaing and L. Gérin. Efficient estimation of the cardinality of large data sets. In *Proceedings of the 4th Colloquium on Mathematics and Computer Science*, volume AG of *Discrete Mathematics & Theoretical Computer Science Proceedings*, pages 419–422, 2006.

[7] P. Clifford and I. A. Cosma. A statistical analysis of probabilistic counting algorithms. *Scandinavian Journal of Statistics*, 2011.

[8] C. Clifton, E. Housman, and A. Rosenthal. Experience with a combined approach to attribute-matching across heterogeneous databases. In *DS-7*, pages 428–451, 1997.

[9] E. Cohen and H. Kaplan. Tighter estimation using bottom k sketches. *PVLDB*, 1(1):213–224, 2008.

[10] R. Cohen, A. Yehezkel, and L. Katzir. A Unified Scheme for Generalizing Cardinality Estimators to Sum Aggregation. *Inf. Process. Lett.*, 2014.

[11] T. Dasu, T. Johnson, S. Muthukrishnan, and V. Shkapenyuk. Mining database structure; or, how to build a data quality browser. In *SIGMOD Conference*, pages 240–251, 2002.

[12] P. Flajolet, É. Fusy, O. Gandouet, and F. Meunier. Hyperloglog: the analysis of a near-optimal cardinality estimation algorithm. In *Analysis of Algorithms (AofA) 2007*. DMTCS.

[13] P. Flajolet and G. N. Martin. Probabilistic counting algorithms for data base applications. *J. Comput. Syst. Sci.*, 31:182–209, Sep. 1985.

[14] É. Fusy and F. Giroire. Estimating the number of active flows in a data stream over a sliding window. In D. Panario and R. Sedgewick, editors, *ANALCO*, pages 223–231. SIAM, 2007.

[15] S. Ganguly, M. N. Garofalakis, R. Rastogi, and K. K. Sabnani. Streaming algorithms for robust, real-time detection of ddos attacks. In *ICDCS*, page 4. IEEE Computer Society, 2007.

[16] F. Giroire. Order statistics and estimating cardinalities of massive data sets. *Discrete Applied Mathematics*, 157:406–427, 2009.

[17] M. A. Hernández and S. J. Stolfo. Real-world data is dirty: Data cleansing and the merge/purge problem. *Data Min. Knowl. Discov.*, 2(1):9–37, 1998.

[18] H. Köhler. Estimating set intersection using small samples. In *ACSC*, pages 71–78, 2010.

[19] K. Krishnamoorthy. *Handbook of Statistical Distributions with Applications*. Chapman & Hall/CRC Press, Boca Raton, FL, 2006.

[20] J. Lumbroso. An optimal cardinality estimation algorithm based on order statistics and its full analysis. In *Analysis of Algorithms (AofA) 2010*. DMTCS.

[21] A. Metwally, D. Agrawal, and A. E. Abbadi. Why go logarithmic if we can go linear?: Towards effective distinct counting of search traffic. In *Proceedings of the 11th International Conference on Extending Database Technology: Advances in Database Technology*, EDBT '08, pages 618–629.

[22] A. E. Monge. Matching algorithms within a duplicate detection system. *IEEE Data Eng. Bull.*, 23(4):14–20, 2000.

[23] E. Rahm and H. H. Do. Data cleaning: Problems and current approaches. *IEEE Data Eng. Bull.*, 23(4):3–13, 2000.

[24] J. Shao. *Mathematical Statistics*. Springer, 2003.

Appendix

Lemma 7
The covariance of $\widehat{|A|}$ *and* $\widehat{|B|}$ *satisfies* $\mathrm{Cov}\left[\widehat{a},\widehat{b}\right] = \frac{n \cdot a \cdot b}{m \cdot u}$, *where* m *is the number of hash functions used for the estimation of* $\widehat{a}$ *and* $\widehat{b}$.

Proof:
Denote x_A and x_B as the maximal hash values for A and B respectively. Using our notation for I_1, I_2 and I_3 from Section 3.1, and applying the linearity of expectation, we obtain:

$$\begin{aligned}\mathbb{E}[x_A \cdot x_B] &= \mathbb{E}[(x_A \cdot x_B) \cdot (I_1 + I_2 + I_3)] \\ &= \mathbb{E}[(x_A \cdot x_B) \cdot I_1] + \mathbb{E}[(x_A \cdot x_B) \cdot I_2] + \mathbb{E}[(x_A \cdot x_B) \cdot I_3]\,. \end{aligned} \tag{19}$$

We consider each term separately:

1) $x_A = x_B$**:** From Eq. (1),

$$\mathbb{E}[(x_A \cdot x_B) \cdot I_1] = \int_{x_A = x_B} n \cdot x_A^{u+1}\, dx = \frac{n}{u+2}.$$

2) $x_A < x_B$**:** From Eq. (2),

$$\mathbb{E}[(x_A \cdot x_B) \cdot I_2] = \iint_{x_A > x_B} \beta \cdot x_B^{\beta} \cdot a \cdot x_A^{a}\, dx_A\, dx_B = \int_0^1 \beta \cdot x_B^{\beta}\, dx_B \cdot \int_0^{x_B} a \cdot x_A^{a}\, dx_A = \\ = \frac{a}{a+1} \cdot \frac{\beta}{u+2}.$$

3) $x_B > x_A$**:**

This case is symmetrical to the second case. We get that

$$\mathbb{E}[(x_A \cdot x_B) \cdot I_3] = \frac{b}{b+1} \cdot \frac{\alpha}{u+2}.$$

Substituting the three terms into Eq. (19) yields that

$$\mathbb{E}[x_A \cdot x_B] = \frac{n}{u+2} + \frac{a}{a+1} \cdot \frac{\beta}{u+2} + \frac{b}{b+1} \cdot \frac{\alpha}{u+2} = \\ = \frac{ab \cdot (u+2) + n}{(u+2)(a+1)(b+1)}. \tag{20}$$

Using the covariance definition, we obtain:

$$\mathrm{Cov}\,[x_A, x_B] = \mathbb{E}[x_A \cdot x_B] - \mathbb{E}[x_A] \cdot \mathbb{E}[x_B] = \frac{ab \cdot (u+2) + n}{(u+2)(a+1)(b+1)} - \frac{a}{a+1} \cdot \frac{b}{b+1} = \\ = \frac{n}{(u+2)(a+1)(b+1)}. \tag{21}$$

The first equality is due to the covariance definition, and the second equality is due to Eq. (20) and Lemma 5(b).

Eq. (21) states the covariance for one hash function. We can generalize it for all hash functions. Let x_A^i and x_B^i be the maximal hash values for the ith hash function, for A and B respectively. Then

$$\mathrm{Cov}\left[x_A^i, x_B^j\right] = \begin{cases} 0 & i \neq j \\ \frac{n}{(u+2)(a+1)(b+1)} & i = j\,. \end{cases} \tag{22}$$

We are now ready to compute $\mathrm{Cov}\left[\widehat{a}, \widehat{b}\right]$. To this end, we first need to choose the cardinality estimator that we will use to estimate $\widehat{a}$ and $\widehat{b}$. For simplicity, we use the estimator from [7]:

$$\widehat{a} = \frac{m}{\sum_{i=1}^{m} (1 - x_A^i)}, \tag{23}$$

where m hash functions are used (symmetrically for $\widehat{b}$). Denoting $X = \sum (1 - x_A^i)$, we can rewrite the estimator as $\widehat{a} = \frac{m}{X}$ (Y is symmetrically defined for $\widehat{b}$). For each i, x_A^i is the maximum of a uniformly distributed variables, and thus

$$\mathbb{E}\left[x_A^i\right] = \frac{a}{a+1} \quad \text{, and} \quad \text{Var}\left[x_A^i\right] = \frac{a}{(a+1)^2(a+2)}. \tag{24}$$

Both equalities follow the beta distribution of x_A^i (Lemma 5(b)). From Eq. (24) it follows that

$$\mathbb{E}[X] = \mathbb{E}\left[\sum (1 - x_A^i)\right] = m \cdot (1 - \mathbb{E}\left[x_A^i\right]) = m \cdot \frac{1}{a+1} = \frac{m}{a+1}, \tag{25}$$

and the variance

$$\text{Var}[X] = \text{Var}\left[\sum (1 - x_A^i)\right] = m \cdot \text{Var}\left[x_A^i\right] = \frac{ma}{(a+1)^2(a+2)}. \tag{26}$$

From Eqs. (25) and (26) we can conclude that $X \to \mathcal{N}\left(\frac{m}{a+1}, \frac{ma}{(a+1)^2(a+2)}\right)$, and symmetrically for Y. Recall that the estimators can be rewritten as $\widehat{a} = \frac{m}{X}$ (symmetrically for $\widehat{b}$). Applying the Delta Method (multivariate version) for the estimator's vector $(\widehat{a}, \widehat{b}) = (m/X, m/Y)$ yields

$$\begin{pmatrix} \widehat{a} \\ \widehat{b} \end{pmatrix} = \begin{pmatrix} m/X \\ m/Y \end{pmatrix} \to \mathcal{N}\left(\begin{pmatrix} a \\ b \end{pmatrix}, g' \cdot \Sigma \cdot g'^T\right),$$

where $g(X, Y) = (m/X, m/Y)$ is the function used in the Delta Method, g' is its partial-derivatives matrix and Σ is the variance matrix:

$$\Sigma = \begin{pmatrix} \text{Var}[X] & \text{Cov}[X, Y] \\ \text{Cov}[X, Y] & \text{Var}[Y] \end{pmatrix}.$$

Using covariance properties and Eq. (22), we obtain:

$$\text{Cov}[X, Y] = \text{Cov}\left[\sum 1 - x_A^i, \sum 1 - x_A^j\right] = \sum_{i=1}^{m} \text{Cov}\left[x_A^i, x_B^i\right] = \frac{mn}{(u+2)(a+1)(b+1)}. \tag{27}$$

Substituting the terms from Eqs. (26) and (27) in Σ yields

$$\Sigma = \begin{pmatrix} \frac{ma}{(a+1)^2(a+2)} & \frac{mn}{(a+1)(b+1)(u+2)} \\ \frac{mn}{(a+1)(b+1)(u+2)} & \frac{mb}{(b+1)^2(b+2)} \end{pmatrix}.$$

Computing $g' \cdot \Sigma \cdot g'^T$ yields the final distribution of the estimators:

$$\begin{pmatrix} \widehat{a} \\ \widehat{b} \end{pmatrix} \to \mathcal{N}\left(\begin{pmatrix} a \\ b \end{pmatrix}, \begin{pmatrix} \frac{a(a+1)^2}{m(a+2)} & \frac{n(a+1)(b+1)}{m(u+2)} \\ \frac{n(a+1)(b+1)}{m(u+2)} & \frac{b(b+1)^2}{m(b+2)} \end{pmatrix}\right).$$

Finally, according to the Delta Method, $\text{Cov}\left[\widehat{a}, \widehat{b}\right]$ is the term in place $[1, 2]$ in the matrix.

Lemma 8
The covariance of $\widehat{|A|}$ *and* $\widehat{|A \cup B|}$ *satisfies* $\mathrm{Cov}\left[\widehat{a}, \widehat{u}\right] = \frac{a^2}{m} + \frac{n \cdot a \cdot b}{m \cdot u} - \frac{u \cdot n \cdot a^2(b^2+\alpha\beta)}{m \cdot Z}$, *where* m *is the number of storage units, and* Z *satisfies:*

$$Z = \alpha \cdot n(a^2 + \alpha\beta) + \beta \cdot n(b^2 + \alpha\beta) + (a^2 + \alpha\beta)(b^2 + \alpha\beta).$$

Proof:
According to Fisher information matrix properties, $\widehat{\mathrm{Cov}\left[\widehat{a}, \widehat{n}\right]} = (\mathbb{F}^{-1})_{1,3}$ is a Maximum Likelihood estimator for $\mathrm{Cov}\left[\widehat{a}, \widehat{n}\right]$, where $(\mathbb{F}^{-1})_{1,3}$ is the term in place $[1, 3]$ in the inverse Fisher information matrix. Computing the term $(\mathbb{F}^{-1})_{1,3}$ yields that

$$\mathrm{Cov}\left[\widehat{a}, \widehat{n}\right] = \frac{u \cdot n \cdot a^2(b^2 + \alpha\beta)}{m \cdot Z}. \tag{28}$$

Therefore,

$$\begin{aligned}\mathrm{Cov}\left[\widehat{a}, \widehat{u}\right] &= \mathrm{Cov}\left[\widehat{a}, \widehat{a} + \widehat{b} - \widehat{n}\right] = \mathrm{Cov}\left[\widehat{a}, \widehat{a}\right] + \mathrm{Cov}\left[\widehat{a}, \widehat{b}\right] - \mathrm{Cov}\left[\widehat{a}, \widehat{n}\right] = \\ &= \frac{a^2}{m} + \frac{n \cdot a \cdot b}{m \cdot u} - \frac{u \cdot n \cdot a^2(b^2 + \alpha\beta)}{m \cdot Z}.\end{aligned}$$

The first equality is due to the inclusion-exclusion principle, and the second is due to covariance properties. The third equality is due to covariance properties, Lemma 2, Lemma 7 and Eq. (28). Similarly, we can obtain the covariance of $\widehat{|B|}$ and $\widehat{|A \cup B|}$:

$$\mathrm{Cov}\left[\widehat{b}, \widehat{u}\right] = \frac{b^2}{m} + \frac{n \cdot a \cdot b}{m \cdot u} - \frac{u \cdot n \cdot b^2(a^2 + \alpha\beta)}{m \cdot Z}.$$